

Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies*

Liran Einav, Amy Finkelstein, and Pietro Tebaldi[†]

June 2024

Abstract: Health insurance is increasingly provided through managed competition, in which two key market design instruments are subsidies for consumers and risk adjustment for insurers. Although typically analyzed in isolation, we illustrate through a stylized model that subsidies offer two key advantages over risk adjustment in markets with adverse selection. First, they provide greater flexibility in tailoring insurance premiums to buyers with different willingness to pay. Second, under imperfect competition, they produce equilibria with lower markups and greater enrollment. We quantitatively assess these effects using demand and cost estimates from the first four years of the California health insurance marketplace regulated by the Affordable Care Act. Holding government spending fixed at approximately the levels in our setting, we estimate that subsidies can increase enrollment by 16 percentage points (76%) over risk adjustment alone, while making all consumers weakly better off.

JEL classification numbers: G22, G28, H51, I13

Keywords: Health Insurance, Market Design, Managed Competition, Insurance Exchanges

*Einav and Finkelstein gratefully acknowledge support from the Sloan Foundation and from the Laura and John Arnold Foundation. Tebaldi acknowledges support from the Becker Friedman Institute. We thank Ben Handel, Mike Whinston, and many seminar participants for helpful comments.

[†]Einav: Department of Economics, Stanford University, and NBER, leinav@stanford.edu; Finkelstein: Department of Economics, MIT, and NBER, afink@mit.edu; Tebaldi: Department of Economics, Columbia University, and NBER, p.tebaldi@columbia.edu.

1 Introduction

In countries around the world, social insurance – particularly pensions and healthcare – represents a large and growing fraction of public expenditures. Historically, social insurance took the form of direct public provision of insurance. Increasingly, however, social insurance involves public regulation of privately provided products. In the context of health insurance in the United States, examples include the approximately one-half (and growing) share of Medicare provided by private firms through the Medicare Advantage program, the 2006 introduction of Medicare Part D for prescription drug coverage provided by regulated private insurers, and the regulated private health insurance marketplaces created by the 2010 Affordable Care Act (ACA). The trend toward regulated competition in health insurance is not limited to the United States; for example, in the Netherlands, Switzerland, and Chile, significant components of their universal coverage are now offered via insurance exchanges similar to those created by the ACA.

In all of these cases, the market sponsor – e.g., a government – sets the rules, and private firms compete within the rules to attract enrollees. Once the sponsor has defined the set of insurance products that can or must be offered, two critical market design decisions remain: premium subsidies for consumers and risk adjustment for insurers. Subsidies are usually viewed as the instrument by which premiums are made affordable to consumers; they are, therefore, often linked to the income of potential buyers. By contrast, risk adjustment systems, which compensate participating insurers for enrolling higher-cost buyers, are typically viewed as a way to reduce concerns about adverse selection and insurer risk skimming.

Perhaps as a result, policy discussions and academic analyses predominantly study subsidies and risk adjustment in isolation, as two separate and unrelated objects. For example, recent work has focused on the impacts of risk adjustment on cream skimming while holding subsidies fixed (e.g., [McWilliams, Hsu, and Newhouse, 2012](#); [Brown, Duggan, Kuziemko, and Woolston, 2014](#); [Geruso and Layton, 2020](#)), or on the effect of premium subsidies on enrollment by low-income individuals while holding risk adjustment fixed (e.g., [Frea, Gruber, and Sommers, 2017](#); [Finkelstein, Hendren, and Shepard, 2019](#); [Tebaldi, Torgovitsky, and Yang, 2023](#); [Tebaldi, forthcoming](#)).

In this paper, we observe that these two instruments – often set by the same entity – naturally interact through their impact on equilibrium allocation. In the stylized price theory framework of [Einav, Finkelstein, and Cullen \(2010\)](#), subsidies are instruments that shift out the demand curve, while risk adjustments are instruments that rotate and shift the cost curve. Given that equilibrium is determined by the intersection of demand and cost, it seems natural to study these two market design features in tandem, and to ask how they may interact and how they substitute or complement each other.

We begin in [Section 2](#) with a standard, stylized model of equilibrium pricing of insurance plans in the presence of adverse selection. In contrast to how subsidies and risk adjustment are typically used in practice, we take a more conceptual approach and allow risk adjustment and subsidies to be functions of the same set of observables. Our main theoretical result is that, for a given level of market sponsor spending, subsidies can achieve higher enrollment and higher consumer surplus than risk adjustment; moreover, they can do so while making each type of consumer (weakly) better off.

This superiority of subsidies over risk adjustment stems from two distinct forces. First, because risk adjustment flattens the cost curve, it increases equilibrium markups, a point noted previously in the literature ([Starc, 2014](#); [Mahoney and Weyl, 2017](#)). We show that, as a result, a (uniform) subsidy can achieve higher coverage and higher consumer surplus at a given level of sponsor spending than risk adjustment can. We refer to this as the “markup effect.” Second, targeted subsidies can reduce inefficiencies arising from adverse selection by targeting different (observable) buyers with different consumer (post-subsidy) premiums, thus incentivizing low-risk types to purchase insurance. We refer to this as the “targeting effect.”

The theory provides qualitative results in a simplified setting. To explore the comparison between subsidies and risk adjustment quantitatively and in a richer setting, the rest of the paper uses the empirical model of [Tebaldi \(forthcoming\)](#) and his estimates of demand and costs from the first four years (2014-2017) of the ACA health insurance marketplace in California. Insurers in the California marketplace offer four standardized coverage options. They must set a uniform premium for each, which is then automatically adjusted by a regulated age factor. During our study period, the California marketplace enrolled more

than a million individuals per year. The vast majority of these enrollees received (age- and income-based) federal subsidies, with approximately 4 billion dollars in annual public expenditures. We consider counterfactual comparisons between risk adjustment and subsidies, while abstracting from many other ACA regulations. Our findings suggest that the qualitative, theoretical observations we point to may be associated with non-trivial quantitative, empirical consequences.

First, to quantify the relevance of our theoretical insights we use our estimates to consider (hypothetical) markets with a single, uniform contract, for which we can compute either competitive or monopoly equilibrium. Holding market sponsor spending fixed, we compare equilibria under different market design regimes: risk adjustment, uniform subsidies, and targeted (non-uniform) subsidies. The enrollment increase in moving from optimal risk adjustment to optimal uniform subsidies reflects the markup effect, while the enrollment increase in moving from uniform to targeted subsidies reflects the targeting effect. Such premium targeting may create a social trade-off between maximizing enrollment and type-specific surplus. Therefore, we also consider potential gains from constrained targeted subsidies that respect a “Pareto” restriction: relative to the uniform subsidy, the constrained targeted subsidies must make all consumers (weakly) better off.

Under perfect competition the markup effect is null, while we find that – at roughly the level of market sponsor expenditures in our setting – the targeting effect leads to enrollment that is 2 percentage points (9%) higher under targeted subsidies relative to risk adjustment and uniform subsidies. Without the “Pareto” constraint on targeted subsidies, enrollment would increase by 4 additional percentage points. For an analogous case of a uniform contract offered by a monopolist, we find that targeted subsidies can increase enrollment by 10 percentage points (80%). The markup effect makes up two-thirds of this difference, while the remaining third is due to the targeting effect.

We then consider the more general (and realistic) case of multi-plan differentiated insurers. In this context, adverse selection operates not only at the extensive margin considered by our theoretical model (by which those buying any coverage at a given premium tend to be riskier) but also at the intensive margin (riskier consumers are more willing to pay to purchase more generous contracts). We estimate that, at roughly the level of market

sponsor expenditures in our setting, targeted subsidies can increase enrollment by about 16 percentage points (76%) while making all consumers weakly better off compared to risk adjustments that perfectly compensate insurers for expected risk. We also find that there is little penalty on enrollment from imposing the “Pareto” constraint on targeted subsidies; relaxing it allows enrollment to increase by only 3 additional percentage points.

In our last set of analyses, we relax our maintained assumption that risk adjustment and subsidies use the same set of observables and that these observables perfectly capture consumer risk. First, we consider more realistic risk scores, which imperfectly predict cost. With a weaker correction to adverse selection, risk adjustment achieves lower enrollment. For example, when risk scores can explain only 30% of the variation in cost, insurance enrollment is about 20% lower than under “perfect” risk adjustment. In addition, we also consider the common situation in which subsidies vary only with income, and are lower for higher-income individuals. Targeting subsidies by income achieves even higher enrollment than targeting on risk since low-income individuals are very sensitive to premium.

The quantitative power of targeted subsidies to achieve welfare gains in our setting is intricately linked to the restriction that insurers cannot price discriminate besides the mandatory age adjustments (so-called “adjusted community rating”). We view some form of community rating as a natural restriction, given that it is widely adopted in regulated health insurance markets. Prior work ([Handel, Hendel, and Whinston, 2015](#)) has illustrated both the costs of community rating in terms of inducing adverse selection as well as its benefits from limiting buyer exposure to reclassification risk (i.e., the risk of subsequent premium changes). Our paper highlights an additional advantage of community rating: it prevents profit-maximizing insurers from undoing the benefits of targeted subsidies via price discrimination. It thus provides the market sponsor with a powerful instrument for increasing insurance enrollment for a given amount of spending.

Our paper is closely related to a large number of papers studying health insurance subsidy design (e.g., [Chan and Gruber, 2010](#); [Decarolis, 2015](#); [Finkelstein et al., 2019](#); [Jaffe and Shepard, 2020](#); [Decarolis, Polyakova, and Ryan, 2020](#); [Tebaldi, forthcoming](#)), risk adjustment (e.g., [Glazer and McGuire, 2000](#); [Wynand, De Ven, and Ellis, 2000](#); [Ellis, 2008](#); [McWilliams et al., 2012](#); [Brown et al., 2014](#); [Layton, McGuire, and Sinaiko, 2016](#); [Einav, Finkelstein,](#)

Kluender, and Schrimpf, 2016; Geruso and Layton, 2020; Saltzman, 2021), the ACA marketplaces (e.g., Abraham, Drake, Sacks, and Simon, 2017; Freaan et al., 2017; Saltzman, 2019; Panhans, 2019; Tebaldi et al., 2023; Dickstein, Ho, and Mark, Forthcoming), and the design of health insurance exchanges more generally (e.g., Handel et al., 2015; Azevedo and Gottlieb, 2017; Curto, Einav, Levin, and Bhattacharya, 2021; Marone and Sabety, 2022; Vatter, 2022). As noted earlier, in all of these papers, subsidies and risk adjustments are treated in isolation, and none of these papers engages in the relationship and tradeoffs between subsidies and risk adjustment that is our focus here.

2 A stylized theoretical framework

2.1 Setting and notation

We consider a stylized setting. There is a single insurance coverage contract, offered by J competing insurers, each indexed by j ; we will relax this assumption in the empirical application. As is often the case in regulated insurance markets, insurers are not allowed to charge different prices to different consumers (beyond any price discrimination that is built into the subsidy design), so each insurer j sets a single price p_j in a Bertrand-Nash Equilibrium. The insurance contract may be horizontally differentiated across insurers (due, e.g., to different provider networks or brand preferences).

Potential buyers are heterogeneous, with each consumer i defined by a triplet (v_i, c_i, w_i) . $v_i = (v_{i1}, \dots, v_{iJ})$ is a vector of consumer i 's willingness to pay for the insurance contracts offered by the different insurers. We denote by $c_i > 0$ the expected cost to the insurer of covering individual i , which, for simplicity, we assume for now to be the same across insurers; we will also relax this assumption in the empirical application. Finally, w_i denotes a vector of observable characteristics, such as age, income, or risk score, which can be used as input to the subsidy or risk adjustment design. We refer to w_i as consumer i 's type.

In this setting, similarly to the framework in Einav et al. (2010), the population is represented by the joint distribution of v_i , c_i , and w_i . This imposes no restrictions on the relationship between preferences, cost, and consumer types. Here, however, we restrict our

analysis to a generic case of adverse selection. In particular, we assume that, for all j and all premium vectors $p = (p_1, \dots, p_J)$, if $AC_j(p)$ is the average cost of individuals covered by contract j , $\partial AC_j(p)/\partial p_j > 0$. As a result, for every j , marginal buyers are cheaper to cover than inframarginal buyers.

A subsidy design is defined as a function $s(w_i)$. If buyer i buys insurance coverage from insurer j , she pays $p_j - s(w_i)$, the market sponsor pays $s(w_i)$, and the insurer's (expected) profits from covering buyer i are $p_j - c_i$.

A risk adjustment design is defined as a function $r(w_i)$. If buyer i is insured, the market sponsor transfers $r(w_i)$ on top of the premium the insurer receives; insurer j 's profits from covering individual i are therefore $p_j - (c_i - r(w_i))$. Although the ACA implemented zero-sum transfers between insurers, we do not require here that risk adjustments be budget neutral; as a result, in principle, risk adjustment payments can result in greater (or lower) overall expenditure by the market sponsor, as is the case, for instance, in Medicare Advantage and in Medicare Part D.

Importantly, throughout the paper we restrict attention to only “ex ante” and “regular” risk adjustment designs, which are the most common in mature markets. By “ex ante” we mean that the risk adjustment function associated with buyer i is known at the time of enrollment, and does not depend on buyer i 's subsequent realized costs or on the realized costs of other buyers in the market.¹ By “regular” we mean that the risk adjustment reduces adverse selection by compensating insurers more generously for covering more risky buyers, or more precisely that $r(w_i) > r(w_k)$ if and only if $\mathbb{E}[c_i|w_i] > \mathbb{E}[c_k|w_k]$.

2.2 Perfect competition

We begin by analyzing the case of perfect competition, which arises in our setting when insurers are homogeneous ($v_{i1} = v_{i2} = \dots = v_{iJ}$, for all i). As a result, as in [Einav et al. \(2010\)](#), the Bertrand-Nash Equilibrium implies that insurers set prices so that price equals average costs, and profits are zero. In such a case, we obtain the following result.

¹In many new markets, it is not uncommon to see such ex-post adjustments that are based on realized costs. As markets mature, however, data availability allows for more accurate risk prediction and a more robust implementation of an ex-ante risk adjustment system.

PROPOSITION 1 *Under perfect competition, for any Nash equilibrium that is achievable with risk adjustment, there exists a subsidy design with no risk adjustment that can achieve the same equilibrium, with the same enrollment for all types and the same total spending by the market sponsor.*

The proof is in the appendix. The intuition is simple and is illustrated in Figure 1. It plots demand and average cost curves as in Einav et al. (2010). The demand curve shows the quantity (or share) of the population who buy insurance at a given price. The marginal cost curve (MC) shows the expected cost for buyers with $v_i = p$, and the average cost curve (AC) shows the expected cost among all consumers who buy insurance at that price. Both curves are downward sloping, indicating the presence of adverse selection: as the price is lowered, the marginal buyer is cheaper to cover than the average existing buyer. Under perfect competition, the equilibrium is given by the intersection of the demand and average cost curve (point A in the top panel of Figure 1).


The top panel of Figure 1 illustrates the impact of risk adjustment. Risk adjustment rotates (flattens) and shifts the average cost curve, leading to a new market equilibrium (point B) with greater insurance coverage and lower insurer prices. In the figure, risk adjustment achieves the efficient outcome, that is the point (point B) at which demand intersects the marginal cost curve.

Proposition 1 implies that this market equilibrium can alternatively be achieved with the same level of market sponsor spending by reverting back to the original average cost curve, and instead using a (uniform) subsidy to accomplish an appropriate parallel shift of the demand curve; this is illustrated in the bottom panel of Figure 1. In the new equilibrium (point C), enrollment is identical to the efficient outcome at point B that was achieved under risk adjustment (top panel), while profits are zero, so market sponsor spending must also be the same.


While the Proposition states that, under perfect competition, any equilibrium that can be implemented via risk adjustment could also be implemented by setting the appropriate level of a uniform subsidy, the converse is not true. That is, it is not the case that any equilibrium under a subsidy design can be implemented using appropriate risk adjustment. To see this, note that under risk adjustment, because insurers must set a single price, all potential

Figure 1: Illustration of Proposition 1

(a) Risk adjustment



(b) Subsidies



Notes: Graphical illustration of Proposition 1. The top panel shows how risk adjustment can move equilibrium from A to the efficient outcome B. The bottom panel shows how subsidies can achieve equilibrium C, where enrollment is identical to B.

customers face the same consumer price, so equilibrium allocations must be monotone in willingness to pay. That is, if $v_i > v_k$ and individual k buys insurance in equilibrium then individual i also does. In contrast, targeted (non-uniform) subsidies could be such that different consumer types face different prices, and if $s(w_k)$ is sufficiently greater than $s(w_i)$ the above monotonicity property could be violated, and the equilibrium could involve some insured individuals having lower willingness to pay than some individuals without insurance. Therefore, targeted subsidies provide additional flexibility to the market sponsor and are able to achieve a greater set of equilibrium allocations than risk adjustment or uniform subsidies.

2.3 Imperfect competition

We now consider a situation in which different buyers may have different valuations for different insurers. In such a situation, insurers have some amount of market power and, in a Bertrand-Nash equilibrium, markups are positive. Here we obtain the following result, where for simplicity we focus on the case of symmetric insurers.

PROPOSITION 2 *Under imperfect competition and adverse selection, for any symmetric Nash equilibrium that is achievable with regular risk adjustment and no subsidies, and in which markups are strictly positive, there is a (uniform) subsidy with no risk adjustment that leads to an equilibrium with the same enrollment for all types, and lower total spending for the market sponsor.*

The formal proof is in the appendix and it is constructive, showing how to calculate the subsidy that leads to a “better” equilibrium for the market sponsor, relative to any risk adjustment. The intuition is similar to the one in [Starc \(2014\)](#) and [Mahoney and Weyl \(2017\)](#), and can be illustrated by examining the following insurer’s first-order condition that must hold in equilibrium:

$$p_j = AC_j(p_j, p_{-j}) - \frac{q_j(p_j, p_{-j})}{\partial q_j(p_j, p_{-j})/\partial p_j} (1 - \partial AC_j(p_j, p_{-j})/\partial p_j), \quad (1)$$

where $q_j(\cdot)$ and $AC_j(\cdot)$ are the residual demand and residual average cost curves faced by the insurer offering plan j . The key object in the first-order condition is $\partial AC_j(\cdot)/\partial p_j$, the

derivative of the $AC_j(\cdot)$ curve with respect to the insurer’s own price. With adverse selection ($\partial AC_j(\cdot)/\partial p_j > 0$), the marginal buyer is cheaper than the average buyer and is therefore relatively attractive to cover, exerting downward pressure on prices and markups. In other words, insurers are more resistant to increasing premiums because the marginal buyers they would lose as a result are relatively cheaper and therefore more attractive to retain.

Regular risk adjustments reduce the difference in costs between the marginal buyer and the average buyer (because the insurers are compensated more generously for covering higher risk individuals), so $\partial AC_j(\cdot)/\partial p_j$ is lower, thus reducing the pressure on prices and leading to greater markups. Uniform subsidies avoid this by not altering the slope of the residual average cost curves. In the perfect competition case the cost reductions implied by risk adjustment are fully passed through to consumers. But when insurers have market power, this pass through is incomplete, and subsidies are therefore a cheaper tool to lower consumer prices (and increase enrollment).

Proposition 2 implies that we can switch from any risk adjustment design to an environment with a uniform subsidy, leading to lower markups, greater coverage, and lower spending by the market sponsor. In addition, the earlier observation about targeted subsidies described in the context of perfect competition remains: risk adjustments (and/or uniform subsidies) imply monotone (in willingness to pay) insurance allocations. Targeted subsidies can relax this, and provide more incentives for the participation of low-cost and/or high-elasticity individuals. Under imperfect competition and adverse selection, both forces put additional downward pressure on prices, and targeted subsidies can lead to equilibria with lower public spending and lower prices for all buyers.²

2.4 Summary

Taken together, these results imply that, for any given level of spending by the market sponsor, subsidies can produce higher coverage (and consumer surplus) than risk adjustment. We call the enrollment difference between risk adjustment and uniform subsidies a “markup effect,” by which subsidies lower insurers’ markups for any coverage level that can be achieved.

²See [Veiga \(2023\)](#) and [Tebaldi \(forthcoming\)](#) for a theoretical and empirical application to targeting subsidies on age. We will provide additional applications in our empirical work below.

We call the additional enrollment gain from moving from uniform to targeted subsidies – and thus allowing the sponsor to provide different subsidies to different types of consumers – a “targeting effect,” by which subsidies can be targeted to provide more incentives for the lower-risk types to enter the insurance pool.

The rest of the paper provides a quantitative assessment of these qualitative results, using data and estimates from the ACA’s health insurance marketplace in California. Parts of our quantitative analysis relax several of the simplifying assumptions of our stylized theoretical setting. In particular, we allow for insurers to offer more than one plan, with vertically (as well as horizontally) differentiated products, and we allow for asymmetric costs across insurers.

3 Empirical setting

The setting, data, empirical specification, and resulting estimates are all taken directly from [Tebaldi \(forthcoming\)](#). We summarize the key features here.

3.1 Setting and data

Setting. Our setting is the first four years (2014-2017) of the California health insurance marketplace (“Covered California”), which was initiated by the Affordable Care Act (ACA). The California marketplace was one of the largest among the fifty states, with more than a million enrolled individuals per year during our sample period.

Covered California partitioned the state into 19 geographic rating regions, each constituting a separate market. Every year, insurers decided whether to participate in the market on a region-by-region basis. We therefore define a market by a region-year combination, and our data covers 76 markets (19 regions over 4 years). There were 3 to 7 insurers participating in each market in our data.

California regulators substantially restricted the scope of insurers’ coverage design. This makes Covered California a useful context in which we can assess the interaction between risk adjustment and subsidies without having to consider other market design elements such as plan features. The regulators required that, within each market, participating insurers

offer all four standardized coverage options. These are labeled by metals – bronze, silver, gold, and platinum – with each metal indicating a different level of coverage generosity, with approximate actuarial values that range from 60% (bronze) to 90% (platinum).³ For silver plans (and only for silver plans), subsidized buyers also received cost-sharing reductions (funded by the government) in addition to premium subsidies. For low-income individuals, these additional subsidies often made silver coverage dominate the corresponding gold coverage, and sometimes even the corresponding platinum coverage.

Despite the standardization of cost-sharing within metal tiers, insurers still differed along two important dimensions. They set different premiums as explained in more detail below, and they offered a different network of medical providers, along with different restrictions on going out of network. These network differences are an important source of heterogeneity across plans that our empirical model will account for.

Premiums were set at the plan level.⁴ Premium setting was subject to regulation, constraining each plan j in a given market to set a single (base) price b_j . This base price was then mapped to the consumer premium p_{ij} , of which $s_{ij} = \min\{s_i, p_{ij}\}$ was paid by the government and $p_{ij} - s_{ij}$ was paid by the individual. These mappings were based on known, pre-specified (by the ACA) formulas, which depended on the individual’s age and household income. Specifically, premiums were the product of the base price and an age factor $p_{ij} = f(\text{age}_i)b_j$,⁵ and s_i was given by $s_i = \max\{0, p_i^* - \bar{p}_i\}$, where p_i^* was equal to the price (p_{ij}) of the second-cheapest silver plan in the market, and $\bar{p}_i = g(\text{income}_i)$ was an increasing function of the individual’s household income.

Data. The data include individual-level enrollment data obtained directly from Covered California and plan-level data on claims from the Center for Medicare and Medicaid Services

³There was a fifth coverage level, “catastrophic coverage,” which offers lower coverage than bronze. It was a high-deductible plan that was only available to individuals who are younger than 35 and who were not eligible for premium subsidies. Since we will limit our empirical analysis to subsidy-eligible individuals, these plans are not relevant for our analysis, and we abstract from them throughout.

⁴There were some cases in which a single insurer offered both HMO and PPO plans in a given metal tier, for a given region and enrollment year. In such cases the insurer would be associated with more than four plans in the (region-year) market.

⁵ $f(\text{age})$ is a monotone function that is increasing from 1 (for 21 years old individuals) to 3 (for 64 years old individuals).

(CMS). The analysis is limited to adults aged 26-64 who are eligible for premium subsidies. This group accounted for 78% of all enrollment in Covered California during our 2014-2017 sample period, for a total of 3.4 million person-years.

For each enrollee the data contain age, household income (as a percentage of the Federal Poverty Level), geographic location, the amount paid by the consumer ($p_{ij} - s_{ij}$), and the plan that was selected. These data are combined with information on premiums (p_{ij}), financial characteristics, and geographic availability for all the plans offered in each market. For most plans, CMS provides the average (ex-post) realized amount of medical claims. Finally, the American Community Survey (ACS) is used to construct a measure of potential buyers across different demographics in each market and the Medical Expenditure Panel Survey (MEPS) is used to calibrate the level of expected medical costs and how they vary by age among insured individuals. Following industry practices, an individual is defined as a potential buyer in a given year (in a given region) if they are either uninsured or purchased insurance in the exchanges or in the individual (non-group) market in the prior year.

Among the more than 12 million potential buyers identified in the ACS, about one third select a plan in Covered California during our study period. On average, these buyers pay almost \$1,500 per year in post-subsidy premiums and receive an average premium subsidy of about \$3,900. After accounting for cost-sharing reductions, the selected plans cover 78% of annual medical spending on average. Most enrollees select either a bronze (24%) or silver (68%) plan. Half of the enrollees select an HMO plan. Except for platinum plans, plan revenues are (on average) substantially above average claims. This is driven by the subsidies; enrollees' payments are (on average) substantially below cost. Consistent with adverse selection, higher-coverage plans are associated with greater total (covered and uncovered) costs.

The average plan in the data enrolls 2,300 individuals. The four largest insurers – Anthem, Blue Shield, HealthNet, and Kaiser – offer the vast majority of the plans observed in the data (889 out of 1,104) and cover 89% of enrollees. This relative dominance mostly reflects regional entry decisions; when present in a market, smaller insurers play an important role, and have an average enrollment of 1,360 individuals in each plan they offer.

3.2 Econometric specification

For each individual i in a given market (rating region and year), we collect their age, household income, and year in the vector $z_i = (age_i, income_i, year_i)$. For each plan j offered in a given market, we observe its base price b_j , its actuarial value AV_j , and a vector of observable characteristics x_j , which consists of insurer indicators (that capture brand reputation and provider network) and an HMO indicator variable (that captures the restriction on using out-of-network providers and, as mentioned earlier, can vary within an insurer-metal tier). For each individual-plan combination, the above observables are then mapped into the implied consumer premium p_{ij} , implied subsidy s_{ij} , and implied actuarial value AV_{ij} .⁶

Following the theoretical framework, the key primitives that govern demand and cost are each individual’s willingness to pay for each plan offered in their market, v_{ij} , and the expected cost to the insurer associated with covering each potential buyer, c_{ij} . We specify each in turn.

Willingness to pay is given by:

$$v_{ij} = \alpha_{z_i}^{-1} (\mu_{z_i} + \beta_i AV_{ij} + \gamma_{z_i} x_j + \xi_{z_{ij}} + \epsilon_{ij}), \quad (2)$$

where the term $\xi_{z_{ij}}$ is an individual-plan characteristic that is known to the insurer, affects individual choice, but is not observed in our data, and ϵ_{ij} is a (standard) iid Type 1 Extreme Value idiosyncratic error term which generates a mixed logit demand specification. In this specification, unobserved heterogeneity across individuals is captured by β_i , the coefficient on the actuarial value of the plan, which is assumed to be lognormally distributed. All other demand parameters vary only with observables.⁷ In practice, the distribution of the ratio $\frac{\beta_i}{\alpha_{z_i}}$ will capture much of the heterogeneity across individuals in willingness to pay for insurance.

⁶More specifically, as described in Section 3.1, the base price is mapped into the premium based on the individual’s age, the subsidy is derived as a function of the individual’s income and the second-cheapest silver plan premium in their market, and the actuarial value varies with individual’s income for silver plans (only for silver plans) due to cost-sharing reductions.

⁷Appendix A of Tebaldi (forthcoming) provides the specific functional forms by which the observables are mapped into each of the other demand parameters.

Insurer j 's expected cost of covering individual i is given by:

$$c_{ij} = AV_j \exp(\rho_i + \phi x_j). \quad (3)$$

Expected cost is thus the expected *total* cost of the individual $\exp(\rho_i + \phi x_j)$, multiplied by the portion of it, AV_j , that is covered by the plan.⁸ Expected total cost in turn is specified as the product of an individual component, $\exp(\rho_i)$, and a plan component that does not vary across individuals. Because $\exp(\rho_i)$ is the only component in c_{ij} that varies across individuals, it will be the key object for risk adjustment in what follows. We construct what we will refer to as a “perfect” risk score by normalizing $\exp(\rho_i)$:

$$w_i = \frac{\exp(\rho_i)}{\mathbb{E} \exp(\rho_i)}, \quad (4)$$

where the expectation is taken over the entire population of potential buyers in the market.

Finally, we specify:

$$\rho_i = \eta^o \text{age}_i + \eta^u \frac{\beta_i}{\alpha_{z_i}}. \quad (5)$$

As mentioned, the ratio β_i/α_{z_i} will capture much of the heterogeneity across individuals in willingness to pay for insurance, so the parameter η^u governs the extent of adverse selection. A positive η^u implies that individuals who tend to have higher willingness to pay for insurance are also those who are more costly to insure.

Estimation and Identification. Demand parameters are estimated using simulated maximum likelihood (where simulations are needed to integrate over the unobserved heterogeneity in β_i), adopting a control function approach to address endogeneity due to correlation between premiums, p_{ij} , and unobserved plan characteristics, $\xi_{z_i,j}$. The parameter η^o , which governs the level of total expected medical cost and how it varies by age among insured individuals, is calibrated directly from data from the MEPS. The remaining cost parameters

⁸Note that this is AV_j rather than AV_{ij} because the cost-sharing subsidies, which are the source of variation in actuarial value across individuals within a plan, are paid by the government and not by the insurer.

(η^u and ϕ) are obtained via simulated method of moments, minimizing the sum of squared differences between observed and model-predicted average claims.

Identification of the demand parameters in equation (2) relies on two sources of variation. First, cost-sharing reductions for enrollees in the silver plans generate sharp discontinuities in AV_{ij} at three income cutoffs (150%, 200%, and 250% of the Federal Poverty Level). Since other characteristics are smooth and continuous in income, differences in choice shares around the cutoffs inform our estimates of willingness to pay for insurance generosity.⁹ Second, because a single pricing decision deterministically affects the premiums paid by individuals of different ages, insurers are more likely to set a higher base price, b_j , in markets with a relatively older population. Since we control for age directly when estimating demand through the μ_{z_i} term, this regulation generates a useful “Waldfoegel-style” instrumental variable (Waldfoegel, 2003): conditional on her age, individual i is more likely to face higher premiums in markets with a higher share of older potential buyers, thus generating identifying variation in prices, which is (arguably) orthogonal to the unobserved individual-plan characteristics in the demand equation, $\xi_{z_i,j}$.

Turning to the cost parameters in equations (3) and (5), plan effects (ϕ) are identified by the correlation between average claims and plan characteristics. The parameter η^u which, as discussed, governs adverse selection is identified by the correlation between residual variation in average claims and the composition of enrollment in terms of β_i/α_{z_i} , identified along with demand.

3.3 Estimates of model primitives

The model and its estimates deliver the primitives highlighted in the theoretical framework, namely for each potential buyer, the joint distribution of (v_i, c_i, w_i) defined in Section 2: individual i 's willingness to pay for each plan j offered in the market, $v_i = (v_{i1}, \dots, v_{iJ})$, the cost to the insurer from covering individual i by each plan j , $c_i = (c_{i1}, \dots, c_{iJ})$, and individual

⁹This research design is similar in spirit to Lavetti, DeLeire, and Ziebarth (2023) who use the income-based discontinuities in cost sharing reductions in silver plans to study the impact of cost-sharing on total medical spending.

Table 1: Summary of Model Estimates

	Mean	Std. Dev.	p10	Median	p90
WTP for 10 pp increase in AV (β/α) (\$US)	426.8	371.2	108.8	320.9	867.5
Prob. of enrollment if all $p = \$1,000$	0.470	0.349	0.055	0.392	0.989
Drop in enrollment if $p = \$1,000 \rightarrow \$1,120$	0.084	0.058	0.002	0.089	0.160
Expected cost for insurer c_{ij} (\$US)	2907	11029	916	2021	5569
Risk score w_i	1	3.79	0.33	0.70	1.90

Notes: The table presents summary statistics of key model estimates across the individuals in our sample. The first three rows show the distribution of key demand estimates; specifically, they show willingness to pay for a 10 percentage points increase in the actuarial value of the coverage, AV (this is driven by the estimate of the ratio β_i/α_{z_i}), the probability of enrolling if all plans' annual consumer premiums were \$1,000, and the percentage drop in this probability if annual consumer premiums increased from \$1,000 to \$1,120. The fourth row summarizes expected cost for the insurer across individuals given their chosen plans (c_{ij}), and the last row shows the risk score (w_i) as defined in equation (4).

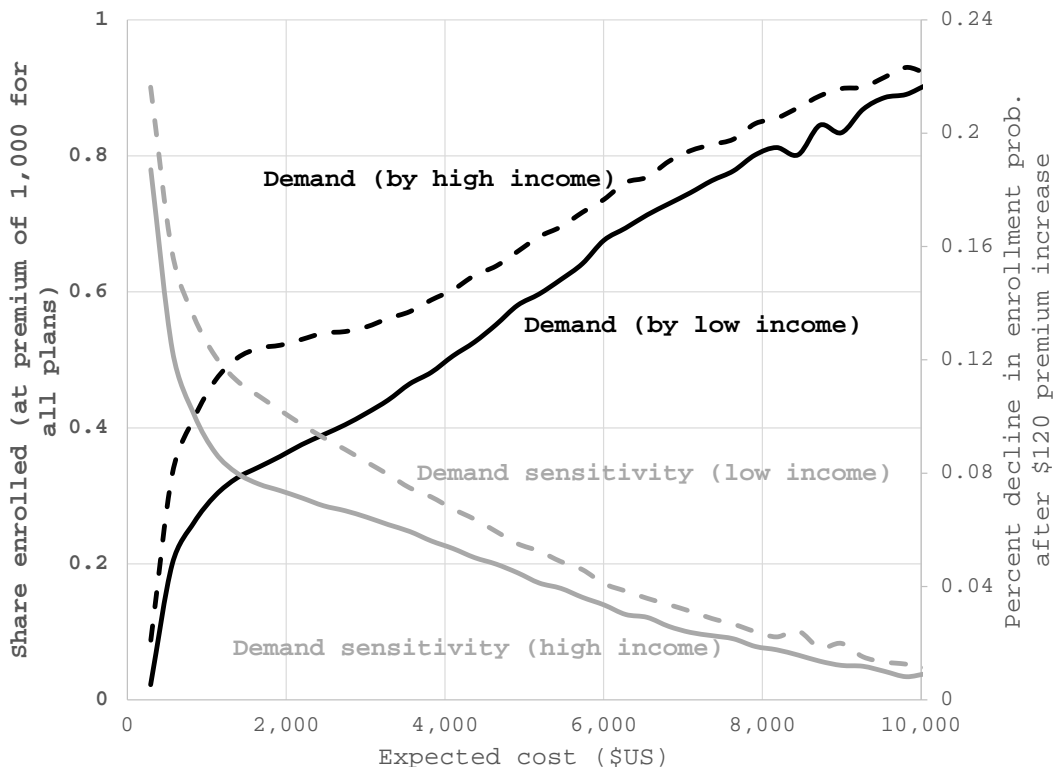
i 's risk score, w_i .¹⁰ These are the primitives that we will use in the next section to generate the counterfactual exercises motivated by the theory laid out in Section 2. Here we provide a short summary that highlights some of the key moments of their joint distribution.

Table 1 summarizes the marginal distributions of v_i , c_i , and w_i . The first row considers the willingness to pay for actuarial value (AV), as measured by the ratio β_i/α_{z_i} . The average actuarial value in our sample is about 80%. We find that individuals are, on average, willing to pay an additional \$427 per year to increase actuarial value by 10 percentage points. This willingness to pay is highly heterogeneous; it is less than \$110 for 10% of the population, but higher than \$870 for another 10%. The next two rows show that this heterogeneity in willingness to pay for coverage translates into substantial heterogeneity in the probability of enrolling in the marketplace. It shows that if all plans' (annual) consumer premiums were set to \$1,000 (about two-thirds the average consumer premium in our sample), enrollment probabilities across individuals would vary a lot (with a median of 0.39), and that if consumer premiums were to increase by \$120 per year (from this \$1,000 baseline), 8% of individuals would drop out of the marketplace.

The penultimate row shows that expected annual costs to the insurer (c_{ij}) are estimated to be \$2,900 for the average individual in our sample, with a large standard deviation of

¹⁰For now, we consider the "perfect" risk score defined in equation (4). At the end of the next section, we will consider alternative, imperfect risk scores.

Figure 2: Expected cost, willingness-to-pay, and demand sensitivity



Notes: The figure summarizes aspects of the joint distribution of demand and cost in our sample. It plots annual expected healthcare spending on the horizontal axis for an individual enrolled in the modal silver plan. The left vertical axis plots the share enrolled in the market when the annual consumer premium is \$1,000 for all plans. The right vertical axis plots the percent decline in the overall probability of enrollment in response to a \$120 per year increase in the consumer premium. We distinguish between individuals with income above and below 250% of the Federal Poverty Line.

\$11,000. These cost are the product of the individual’s expected annual medical spending and their chosen plan. Based on data from the Medical Expenditure Panel Survey, we estimated average projected annual medical spending for the demographic group we study to be just over \$4,100, but there is of course substantial dispersion around that mean with a standard deviation of \$12,000. The last row of the table summarizes the individual’s risk score w_i , defined in equation (4), which is invariant across plans. Individuals at the 90th percentile of the risk distribution have expected medical expenditures that are almost six times as large as those at the 10th percentile.

The quantitative results of our counterfactual exercises will depend not only on the marginal distributions of the triplet (v_i, c_i, w_i) – which are summarized in Table 1 – but also on their joint distribution. Figure 2 therefore illustrates how demand and demand sensitivity vary with expected cost. We plot expected annual cost for the insurer when the individual

is enrolled in the modal silver plan on the horizontal axis. On the left vertical axis we show the average probability of enrolling in the market if all consumer premiums for all plans are \$1,000; we show this separately for individuals with income above or below 250% of the Federal Poverty Level. On the right vertical axis we show the decline in the probability of enrollment when premiums go up, separately for the same income groups.

Figure 2 shows the presence of adverse selection in our empirical context: within a given income group, higher-cost individuals are, on average, significantly more willing to pay for insurance. It also shows that at a given cost, higher income individuals are more likely to purchase coverage and are less sensitive to premiums increases, as are higher cost individuals. For example, we estimate that, at consumer premiums of \$1,000, about 50% of the high income group with an expected cost of \$2,000 would purchase coverage. This share is about 15 percentage points lower among those with income lower than 250% of the Federal Poverty. Moreover, if premiums were to increase by \$120 per year, there would be a decline of about 10% in the enrollment probability of low-income individuals with an expected cost of \$2,000, and a smaller decline (of about 7%) in the enrollment probability of high-income individuals. In comparison, at the same level of premiums, about 90% of those with an expected cost of \$10,000 would purchase coverage, and income differences in price sensitivity are smaller. For this group, a \$120 premium increase would push less than 1% of enrollees out of the marketplace.

From Figure 2 we can also see that it is possible to “rank” individuals in such a way that those with lower expected cost are less likely to be enrolled and more sensitive to premium changes, and vice versa. As shown in Section 2, this adverse selection is sufficient for subsidies to perform better than risk adjustment. We now turn to a quantitative evaluation of alternative market designs in this setting.

4 Comparisons of alternative market designs

We use the California marketplace and the estimates of demand and cost described in the last section in order to solve for market equilibria under alternative, counterfactual risk adjustment or subsidy schemes. Naturally, we have many degrees of freedom in the way

we perform these exercises. We do not view these exercises as providing a precise policy prescription. Rather, they are illustrative of the range of plausible magnitudes associated with the theoretical mechanisms highlighted in Section 2, as well as the robustness of these mechanisms to relaxing some of the assumptions in the theoretical propositions.

4.1 Setup

We assume that the market sponsor’s objective is to maximize overall enrollment subject to the budget constraint by picking from a class of alternative market design regimes we consider. Motivated by the theoretical results in Section 2, we consider the following four market design regimes:

- A *risk adjustment* regime, in which the market sponsor chooses the risk adjustment payment function $r(w_i) = \mu \cdot w_i$ with $\mu > 0$. That is, insurers are paid an amount for each enrollee that is proportional to the enrollee’s risk score. Higher values of μ imply more aggressive risk adjustment and a flatter average cost curve.
- A *uniform subsidy* regime, in which the market sponsor chooses a single number s , which is the uniform subsidy that potential buyers can use toward premium payments.
- A *targeted subsidy* regime, in which the market sponsor chooses a function $s(w_i)$, which represents the subsidy amount that a potential buyer of type w_i can use toward premium payment. In particular, we consider functions of the form $s(w_i) = \lambda_0(1 - \lambda_1(1 - 1/w_i))$. Here, λ_0 controls the “level” of subsidies, while λ_1 controls its “targeting.” Setting $\lambda_1 = 0$ corresponds to the case of uniform subsidies, while values of $\lambda_1 > 0$ provide higher subsidies to lower-cost individuals and lower subsidies for higher-cost individuals. The larger the value of λ_1 , the stronger the targeting of subsidies toward lower-cost potential buyers.
- A *constrained targeted subsidy* regime, which is the same as the targeted subsidy regime, except that the values of λ_0 and λ_1 are chosen subject to the additional constraint that at the equilibrium outcome, the move from a uniform subsidy to a targeted one does not harm any consumer; that is, no consumer pays a higher premium.

For each of these four market design regimes, we solve for the design parameters that maximize the sponsor’s objective (total enrollment) for a given amount of spending by the market sponsor. Specifically, we compute a Bertrand-Nash equilibrium under a wide range of values of the market design parameters, and then search for the ones that maximize the objective subject to the sponsor’s budget constraint; Appendix B describes the computation in more detail. Thus, for each value of the budget, we obtain four different solutions, one for each market design regime, that are optimal over the range of designs we consider. To facilitate meaningful comparisons across alternatives, we will often compare outcomes across market design regimes holding fixed the total amount of spending by the market sponsor, expressed in dollars per year per potential customer.

We conduct two main types of exercises. First, in Section 4.2 we use the empirical estimates from the California marketplace but simplify the setting to consider only one homogeneous (silver) plan offered competitively or by a monopolist. This simplified setting makes the empirical context match the stylized setting of Section 2; it guarantees that our theoretical predictions hold, and allows us to provide intuition for the key forces that are at play. Second, in Section 4.3 we undo many of the simplifications and approximate more closely the actual setting of the California marketplace by redoing the same exercise considering all available bronze and silver plans.¹¹ This allows us to obtain quantitative estimates that are more meaningful, but with arguably less clear intuition than the simplified setting.


4.2 Perfect competition or monopoly with a homogeneous plan

We begin with a simplified setting in which there is only single coverage option, which we assume to be the most common silver plan offered in each market.

A single market. We first focus on one specific market, the 2014 Los Angeles rating region, to provide empirical illustrations of the theoretical propositions from Section 2. Figure 3 illustrates Proposition 1 by generating the empirical analog to Figure 1. The black lines

¹¹To reduce computation time and limit concerns of equilibrium multiplicity, we omit gold and platinum plans, whose joint market share is quite small (less than 8%).

Figure 3: Empirical Example of Proposition 1: Perfect Competition




Notes: The figure considers a homogeneous silver plan offered by multiple insurers in a perfectly competitive market, using estimates from the Los Angeles rating region in 2014. We plot demand, marginal cost, and average cost without any intervention of the market sponsor, risk adjustment achieving the efficient outcome (which corresponds to $\mu = \$1,000$ with 46,000 enrollees), and a uniform subsidy at the same level of market sponsor spending (which corresponds to a subsidy s of \$2,055).

show demand and cost curves without any intervention by the market sponsor. Demand and marginal cost cross when 46,000 individuals purchase coverage; this would be the efficient outcome. The gray lines illustrate two alternative market designs, each achieving the efficient outcome with a total spending of 96 million dollars. At this level of spending, risk adjustment – which shifts and flattens the average cost curve – leads to one market equilibrium (indicated by the lower red dot), while uniform subsidies – which shift the demand curve – leads to another equilibrium indicated by the second red dot. Both designs achieve the same enrollment (46,000, approximately 9% of potential buyers), with the same market sponsor spending, illustrating Proposition 1. Note that the vertical distance between the two equilibrium outcomes represents the size of the subsidy (\$2,055); consumer premiums are the same under both regimes.


Figure 4 shows what happens with the same two instruments – risk adjustment and

Figure 4: Empirical Example of Proposition 2: Monopoly

(a) Risk Adjustment



(b) Uniform Subsidies



Notes: The figure considers a homogeneous silver plan offered by a monopolist, using estimates from the Los Angeles rating region in 2014. In both panels we plot demand, marginal cost, average cost, and marginal revenue. The black lines illustrate these objects without any intervention by the market sponsor. The gray lines illustrate them under alternative market designs. The top panel considers the case of risk adjustment with $\mu = \$2,170$ to achieve the efficient enrollment level of 46,000 (as in Figure 3). The resulting equilibrium markup is \$2,383 (indicated by the vertical red arrow) and market sponsor spending is 210 million (not shown). In the bottom panel, we follow Proposition 2 and find the enrollment-equivalent uniform subsidy. This subsidy is equal to \$2,657 (not shown); enrollment is the same (by design, as indicated by the red dot) while the markup (indicated by the vertical red arrow) is much lower (\$602) as is market sponsor spending which totals \$124 million dollars (not shown).

uniform subsidy – when this market is now a monopoly. The top panel shows the results with risk adjustment achieving the efficient enrollment level as in the perfectly competitive case in Figure 3. This now requires a higher $\mu = \$2,170$, since the pass-through of sponsor’s payments to consumers is imperfect due to market power. At the equilibrium with this risk adjustment the average markup is \$2,383, and the total spending by the market sponsor to achieve the efficient outcome increases to 210 million dollars, more than twice as large as under perfect competition.

More interestingly for our purposes, in the bottom panel of Figure 4 we follow the construction in Proposition 2 and use a uniform subsidy to achieve the same enrollment. Under the uniform subsidy ($\mu = \$2,657$), the markup is now substantially lower (\$602 compared to \$2,383 under risk adjustment). The uniform subsidy is thus able to achieve the same enrollment as risk adjustment, but at substantially lower spending by the market sponsor: total spending drops by 40% to 124 million dollars.

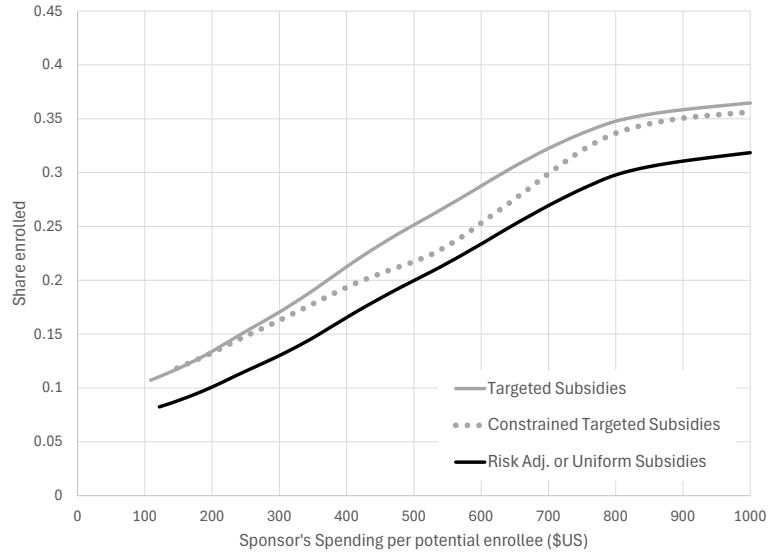
All markets. In Figure 5 we continue to maintain a homogeneous coverage option, but now consider all markets in our data (rather than just one), and a greater range of market designs. We continue to examine two extreme market structures: a monopolistic market and a perfectly competitive market in which multiple insurers offer the same plan, so equilibrium profits are zero. We plot average enrollment shares (across all markets) as a function of the sponsor budget (expressed in dollars per potential buyer).

The top panel of Figure 5 considers the case of perfect competition. Here, there is no markup effect so risk adjustment and uniform subsidies are equivalent (solid black line). In contrast, targeted subsidies generate higher enrollment for any level of spending, with constrained targeted subsidies achieving more than half of the benefits derived from targeting (for most levels of spending). At the average level of sponsor spending under the ACA, which is just over \$600 per potential buyer,¹² the targeting effect leads to enrollment that is 2 percentage points (9%) higher under targeted subsidies relative to risk adjustment and


¹²In 2015, Federal subsidies in ACA marketplaces totaled approximately 31.8 billion dollars (Table 6 in https://aspe.hhs.gov/sites/default/files/migrated_legacy_files//137886/ib_2015mar_enrollment.pdf) while the size of the individual market was approximately 50.6 million individuals (<https://www.kff.org/other/state-indicator/total-population/>), for an average spending per potential buyer of \$630.

Figure 5: Risk Adjustment vs. Subsidies: A homogeneous plan

(a) Perfect Competition



(b) Monopoly



Notes: The figure considers the simple case of a homogeneous silver plan in every market, and plots the share enrolled as a function of market sponsor spending (per potential buyer) across the 76 region-year markets in our data. For “constrained targeted subsidies” we consider the highest level of enrollment achieved by any rotation (λ_1), imposing that all individuals pay weakly lower premiums compared to the case of uniform subsidies. For unconstrained targeted subsidies we consider the highest level of enrollment achieved by any rotation (λ_1), but now without the additional constraint that all individuals pay premiums weakly lower than the uniform case. The figure is constructed by simulating equilibrium for every design and every market separately over a grid of values for μ , λ_0 , and λ_1 . We then interpolate enrollment and spending using local polynomials. Appendix B provides more details.


uniform subsidies. Without the “Pareto” constraint on targeted subsidies, enrollment would increase by 4 additional percentage points.

The bottom panel examines the same four designs under monopoly. Naturally, vis-à-vis perfect competition, enrollment is lower due to market power, which leads to higher equilibrium prices with the insurer capturing some share of the surplus. For example, at any level of sponsor spending, risk adjustment leads to enrollment that is just over half of what enrollment would be under perfect competition. More importantly for our purpose, in the context of monopoly – unlike in the case of perfect competition – the equilibrium under risk adjustment and under uniform subsidies now differ substantially due to the markup effect; this is indicated by the difference in enrollment between risk adjustment and uniform subsidies (holding spending fixed) in the bottom panel of Figure 5. For example, at the average level of sponsor spending under the ACA (\$600 per potential buyer), the share of individuals purchasing coverage under risk adjustment is about 0.13, but it increases by 54%, to 0.2, under uniform subsidies. Targeted subsidies achieve even higher enrollment, bringing the share enrolled up to 0.23, or about an 80% increase relative to risk adjustment, although the additional 3 percentage point enrollment gain over uniform subsidies is smaller than the 6 percentage point enrollment gain from moving from uniform subsidies to targeted subsidies under perfect competition (top panel). Targeting subsidies to lower-risk individuals attracts a larger share of them to the market, which reduces average costs; however, this reduction is partly captured by the monopolist in the form of higher profits, rather than perfectly passed through to lower premiums (which is the case with perfect competition), which is why the targeting effect is smaller under monopoly.

4.3 Oligopoly with differentiated plans

Relative to the simplified setting described in the last section, we now consider the full set of (heterogeneous) insurers that operate in each of the 76 markets in our data. In other words, rather than restrict insurers to offer a homogeneous (silver) plan within the market, we now allow these plans to be differentiated across insurers (e.g. based on provider network and/or insurer brand preferences). In this richer setting that includes vertical differentiation between the bronze and silver plans offered by each insurer, adverse selection now operates on

Figure 6: Risk Adjustment vs. Subsidies: Heterogeneous Multiproduct Oligopoly



Notes: The figure considers differentiated multi-plan insurers, each offering bronze and silver plans, and plots the share enrolled as a function of market sponsor spending (per potential buyer) across the 76 region-year markets in our data. For “constrained targeted subsidies” we consider the highest level of enrollment achieved by any rotation of the subsidy schedule (λ_1), imposing that all individuals pay weakly lower premiums compared to the case of uniform subsidies. For unconstrained targeted subsidies we consider the highest level of enrollment achieved by any rotation (λ_1), but now without the additional constraint that all individuals pay premiums weakly lower than under risk adjustment. The figure is constructed by simulating equilibrium for every design and every market separately over a grid of values for μ , λ_0 , and λ_1 . We then interpolate enrollment and spending using local polynomials. Appendix B provides more details.

two margins (as studied in Geruso, Layton, McCormack, and Shepard (2023)): the extensive margin of whether the individual is covered (as we have considered thus far) and the intensive margin of the amount of coverage (since riskier individuals are also more willing to pay for generous coverage). In this more complicated setting, we often fail to find equilibria for the case of uniform subsidies, which are therefore not considered in this part of our analysis.

Figure 6 illustrates our results graphically, comparing enrollment across the three designs with varying levels of spending by the market sponsor. The ordering between risk adjustment and subsidies highlighted in Section 2 appears to be robust to relaxing the assumptions of symmetric insurers offering only one plan. Even in this richer setting – with horizontal and vertical differences between plans – we find that the magnitudes are far from trivial. For a level of spending similar to the one observed under the ACA – approximately \$600 (see footnote 12) – targeted subsidies achieve a 76% increase in enrollment (from 0.21 to 0.37)

Table 2: Comparing Designs at Sponsor Spending of \$600 per Potential Buyer

	Risk Adjustment	Constrained Targeted Subsidies	Targeted Subsidies
Share enrolled	0.21	0.37	0.40
Share of enrollees with below-median risk score	0.30	0.45	0.43
Premium paid by enrollees with highest risk score (\$US)	1,279	633	1,294
Sponsors’s spending per enrollee (\$US)	3,014	1,691	1,567

Note: The table summarizes equilibrium outcomes for the case of heterogeneous insurers offering bronze and silver plans across different risk adjustment and subsidy designs. Spending by the market sponsor is held fixed at \$600 per potential buyer, the average ACA level (see footnote 12).

relative to risk adjustment, even constraining all consumers to be (weakly) better off. Upon relaxing the requirement of all buyers paying prices weakly lower than under risk adjustment, targeting would achieve even higher levels of aggregate enrollment (0.40).

Figure 6 also indicates that the enrollment difference achieved by constrained targeted subsidies relative to risk adjustment is much smaller at lower levels of spending by the sponsor (and therefore lower enrollment). This is because at lower levels of enrollment, the average cost curve is much steeper and many marginal buyers mostly consist of high-income, high-risk individuals (see Figures 2, 3 and 4). Therefore, at these low level of spending adverse selection – and, more precisely, the derivative of AC with respect to premium, c.f. equation (1) – remains substantial even under risk adjustment, keeping markups low. At the same time, the scope for targeting is also limited unless one were to relax the constraint that all consumers are (weakly) better off and adopt policies that exclude high-risk individuals from the market (gray line).

Table 2 summarizes a number of other outcomes across risk adjustment and subsidy designs holding spending to the average level of \$600 per potential buyer. It shows enrollment rates overall and the share of enrollees with below-median risk score. It also shows the highest consumer premium that anyone pays (across all types, in our analysis this is the premium paid by those with the highest risk score) and the market sponsor spending per enrollee.

In addition to the total enrollment effects, the table provides some insight into what is happening “under the hood.” Under risk adjustment, consumer premiums are the same for everyone, and the average premium paid is equal to \$1,279. Moving to constrained targeted subsidies - in which subsidies are allowed to vary across types but with the constraint that

no consumer pays a higher premium than under risk adjustment - enrollment increases and the proportion of low-risk enrollees is 15 percentage points higher (increasing from 30 to 45%). The effect of increasing the number of low cost individuals in the enrollment pools is substantial: even if different individuals pay different premiums, the highest premium paid by any enrollee to obtain coverage is \$633, approximately half as large as under risk adjustment.


The last column shows that higher risk individuals would be worse off if the constraint on targeted subsidies were relaxed, and the market sponsor could target subsidies to maximize enrollment without requiring that all consumers be weakly better off than under uniform subsidies. Absent this constraint, enrollment would be 3 percentage point higher, but the highest premium in the market now increases from \$633 to \$1,294 for the highest risk individuals.

4.4 More realistic targeting

In all the analyses thus far, we have considered “perfect” observability of types, so that w_i tracks “perfectly” individual expected cost (see equation (4)). We now explore two more realistic situations in which the observables used in risk adjustment and in subsidies are neither perfect nor identical. For targeted subsidies, we now allow λ_0 and λ_1 to vary only with income, and consider the common case in which higher-income individuals receive lower subsidies. For risk scores, we consider risk scores that approximate the current state-of-the-art risk scoring models, which have an R^2 from a regression of medical spending on risk scores that is lower than 0.3 (Kautter et al., 2014; McGuire, Zink, and Rose, 2021), far from the ideal case of $R^2 = 1$ that we have considered so far. Specifically, we replace the perfect risk score from above with a noisy one, which we design to be such that the R^2 of our risk adjustment model is 0.3. Appendix B provides more details.

Figure 7 summarizes the enrollment achieved under these various designs, holding market sponsor spending per potential buyer fixed at \$600 as in Table 2. We normalize enrollment under the perfect risk adjustment with $R^2 = 1$ (see Table 2) to be one. Adding noise to the risk adjustment decreases enrollment by 20%. This is the opposite sign of what would be expected based on the stylized model we presented in Section 2, but as noted, we are now

Figure 7: Shared Enrolled Relative to Risk Adjustment with $R^2 = 1$



Note: The figure summarizes the comparison of enrollment achieved by different designs, holding market sponsor spending fixed at \$600 per potential buyer. Enrollment under “perfect” risk adjustment with $R^2 = 1$ is normalized to one. We then consider noisy risk adjustment with $R^2 = 0.3$. This is constructed by setting $w_i = e^{\rho_i} / \mathbb{E}[e^\rho] + \omega_i$, where ω_i is drawn iid randomly across individuals to target the R-square of a linear regression of cost on w_i to be 0.3. The third bar corresponds to targeted subsidies, already analyzed in Section 4. We then consider subsidies targeted only on income, as common in many markets.

in a richer setting in which the qualitative theoretical results need not hold.

Interestingly, we also find that targeting subsidies on income performs even better than subsidies on risk. Quantitatively, providing higher subsidies to low-income consumers leads to total enrollment that is 14% higher. This is consistent with the patterns illustrated in Figure 2: since high-risk individuals almost always purchase coverage irrespective of their income, allocating the marginal dollar of sponsor spending to low-income individuals amounts to targeting those with low risk and also the highest sensitivity to premiums. This implies a stronger impact on average cost and average premium elasticities, and thus a stronger effect on premiums and enrollment. A comparison between noisy risk adjustment and income-targeted subsidies – the two policies that one may expect to see more often in marketplaces – shows that enrollment under the latter design would 2.5 times as large. This result could be a justification for the widespread use of income-targeted subsidies in the ACA exchanges.

5 Conclusions

Our paper has emphasized the importance of jointly considering subsidies and risk adjustment – two common market design instruments often employed by the same market sponsor – rather than to analyze each in isolation, as is typically done in both academic and health policy circles. Once we recognize that by altering market demand and market costs, respectively, subsidies and risk adjustment jointly interact to determine market equilibrium, the standard practice of thinking about subsidies as a way to achieve “affordability” and risk adjustment as a way to ameliorate adverse selection seems unsatisfactory.

We show theoretically that, at least under very stylized assumptions, subsidies can achieve greater enrollment for a given level of market sponsor spending. Using data and existing estimates from California’s ACA health insurance marketplace during 2014-2017, and relaxing many of the stylized assumptions of the theory, we estimate that, holding sponsor spending fixed at roughly the level of federal subsidy expenditures in this market, subsidies can increase enrollment by about 16 percentage points (76%) compared to risk adjustment that perfectly compensates insurers for risk. Further, this increase in enrollment is achieved while holding all types (weakly) better off compared to the equilibrium with risk adjustment.

A natural question raised by our theoretical and empirical results is why risk adjustment remains a popular market design instrument. One possible economic explanation is that risk adjustment serves other functions beyond its role in pricing that we considered here. In particular, by decreasing the relative profitability to insurers of healthier compared to less healthy enrollees, risk adjustment may be important for reducing insurer incentives to cream-skimming using non-price instruments, such as benefit design or marketing.¹³

There are also potential political economy explanations for the continued use of risk adjustment. For example, while our theoretical and empirical analyses allow risk adjustment and subsidies to condition on the same type space, in practice offering greater subsidies to healthy consumers – as optimal subsidies would often require – may conflict with naive intuition and may be politically difficult. Likewise, insurer profits may be higher under risk adjustment, creating a potential political force in favor of them. In this sense, our results can be thought of as providing a quantitative assessment of the costs of such potential constraints, in the context of California’s ACA health insurance exchange.

More broadly, our intent here is not to prescribe specific market design strategies for health insurance exchanges, but rather to highlight the important sense in which two market design tools are highly related, and to provide some quantitative assessment of the trade-off associated with greater reliance on risk adjustment relative to a richer and more flexible subsidy design.

References

- ABRAHAM, J., C. DRAKE, D. W. SACKS, AND K. SIMON (2017): “Demand for health insurance marketplace plans was highly elastic in 2014–2015,” *Economics Letters*, 159, 69–73.
- AZEVEDO, E. M. AND D. GOTTLIEB (2017): “Perfect competition in markets with adverse selection,” *Econometrica*, 85, 67–105.
- BROWN, J., M. DUGGAN, I. KUZIEMKO, AND W. WOOLSTON (2014): “How does risk

¹³Another potential economic rationale for risk adjustment is that it allows conditioning payments on ex post realized costs, whereas subsidies must be based on ex ante measures. However, ex-post risk adjustment seems suboptimal, as it increases gaming opportunities (see e.g. [Geruso and Layton, 2020](#)), and this may be why we rarely see it in mature markets.

- selection respond to risk adjustment? New evidence from the Medicare Advantage Program,” *American Economic Review*, 104, 3335–3364.
- CHAN, D. AND J. GRUBER (2010): “How sensitive are low income families to health plan prices?” *American Economic Review*, 100, 292–296.
- CURTO, V., L. EINAIV, J. LEVIN, AND J. BHATTACHARYA (2021): “Can health insurance competition work? evidence from medicare advantage,” *Journal of Political Economy*, 129, 570–606.
- DECAROLIS, F. (2015): “Medicare part d: Are insurers gaming the low income subsidy design?” *American Economic Review*, 105, 1547–1580.
- DECAROLIS, F., M. POLYAKOVA, AND S. P. RYAN (2020): “Subsidy design in privately provided social insurance: Lessons from medicare part d,” *Journal of Political Economy*, 128, 1712–1752.
- DICKSTEIN, M. J., K. HO, AND N. D. MARK (Forthcoming): “Market Segmentation and Competition in Health Insurance,” *Journal of Political Economy*.
- EINAIV, L., A. FINKELSTEIN, AND M. R. CULLEN (2010): “Estimating welfare in insurance markets using variation in prices,” *The quarterly journal of economics*, 125, 877–921.
- EINAIV, L., A. FINKELSTEIN, R. KLUENDER, AND P. SCHRIMPF (2016): “Beyond statistics: the economic content of risk scores,” *American Economic Journal: Applied Economics*, 8, 195–224.
- ELLIS, R. P. (2008): “Risk adjustment in health care markets: concepts and applications,” *Paying for Health Care: New Ideas for a Changing Society*. Wiley-VCH publishers Weinheim, Germany.
- FINKELSTEIN, A., N. HENDREN, AND M. SHEPARD (2019): “Subsidizing health insurance for low-income adults: Evidence from Massachusetts,” *American Economic Review*, 109, 1530–1567.
- FREAN, M., J. GRUBER, AND B. D. SOMMERS (2017): “Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act,” *Journal of health economics*, 53, 72–86.
- GERUSO, M. AND T. LAYTON (2020): “Upcoding: evidence from Medicare on squishy risk adjustment,” *Journal of Political Economy*, 128, 984–1026.
- GERUSO, M., T. J. LAYTON, G. MCCORMACK, AND M. SHEPARD (2023): “The two-margin problem in insurance markets,” *Review of Economics and Statistics*, 105, 237–257.
- GLAZER, J. AND T. G. MCGUIRE (2000): “Optimal risk adjustment in markets with adverse selection: an application to managed care,” *American Economic Review*, 90, 1055–1071.

- HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): “Equilibria in health exchanges: Adverse selection versus reclassification risk,” *Econometrica*, 83, 1261–1313.
- JAFFE, S. AND M. SHEPARD (2020): “Price-linked subsidies and imperfect competition in health insurance,” *American Economic Journal: Economic Policy*, 12, 279–311.
- KAUTTER, J., G. C. POPE, M. INGBER, S. FREEMAN, L. PATTERSON, M. COHEN, AND P. KEENAN (2014): “The HHS-HCC risk adjustment model for individual and small group markets under the Affordable Care Act,” *Medicare & Medicaid research review*, 4.
- LAVETTI, K., T. DELEIRE, AND N. R. ZIEBARTH (2023): “How do low-income enrollees in the Affordable Care Act marketplaces respond to cost-sharing?” *Journal of Risk and Insurance*, 90, 155–183.
- LAYTON, T. J., T. G. MCGUIRE, AND A. D. SINAIKO (2016): “Risk corridors and reinsurance in health insurance marketplaces: insurance for insurers,” *American journal of health economics*, 2, 66–95.
- MAHONEY, N. AND E. G. WEYL (2017): “Imperfect Competition in Selection Markets,” *The Review of Economics and Statistics*, 99, 637–651.
- MARONE, V. R. AND A. SABETY (2022): “When Should There Be Vertical Choice in Health Insurance Markets?” *American Economic Review*, 112, 304–342.
- MCGUIRE, T. G., A. L. ZINK, AND S. ROSE (2021): “Improving the performance of risk adjustment systems: constrained regressions, reinsurance, and variable selection,” *American Journal of Health Economics*, 7, 497–521.
- MCWILLIAMS, J. M., J. HSU, AND J. P. NEWHOUSE (2012): “New risk-adjustment system was associated with reduced favorable selection in Medicare Advantage,” *Health Affairs*, 31, 2630–2640.
- PANHANS, M. (2019): “Adverse selection in ACA exchange markets: Evidence from Colorado,” *American Economic Journal: Applied Economics*, 11, 1–36.
- SALTZMAN, E. (2019): “Demand for health insurance: Evidence from the California and Washington ACA exchanges,” *Journal of Health Economics*, 63, 197–222.
- (2021): “Managing adverse selection: underinsurance versus underenrollment,” *The RAND Journal of Economics*, 52, 359–381.
- STARC, A. (2014): “Insurer pricing and consumer welfare: Evidence from medigap,” *The RAND Journal of Economics*, 45, 198–220.
- TEBALDI, P. (forthcoming): “Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca,” *Review of Economic Studies*.
- TEBALDI, P., A. TORGOVITSKY, AND H. YANG (2023): “Nonparametric estimates of demand in the california health insurance exchange,” *Econometrica*, 91, 107–146.

- VATTER, B. (2022): “Quality disclosure and regulation: Scoring design in medicare advantage,” *Available at SSRN 4250361*.
- VEIGA, A. (2023): “Price Discrimination in Selection Markets,” *Review of Economics and Statistics*, 1–45.
- WALDFOGEL, J. (2003): “Preference externalities: An empirical study of who benefits whom in differentiated-product markets,” *The Rand Journal of Economics*, 34, 557.
- WYNAND, P., V. DE VEN, AND R. P. ELLIS (2000): “Risk adjustment in competitive health plan markets,” in *Handbook of health economics*, Elsevier, vol. 1, 755–845.

APPENDIX

A Theoretical Appendix

Definitions. Given the primitives introduced in Section 2, the average cost function for contract j is

$$AC_j(p) = \frac{C_j(p)}{q_j(p)} = \frac{\int \mathbf{1} \{v_{ij} - p_j \geq v_{ik} - p_k \text{ for all } k, \text{ and } v_{ij} - p_j \geq 0\} c_i dF(i)}{\int \mathbf{1} \{v_{ij} - p_j \geq v_{ik} - p_k \text{ for all } k, \text{ and } v_{ij} - p_j \geq 0\} dF(i)},$$

where $C_j(\cdot)$ and $q_j(\cdot)$ are, respectively, the total cost and total enrollment functions.

Under a risk adjustment $r(\cdot)$ average cost becomes

$$AC_j^r(p) = \frac{C_j(p)}{q_j(p)} = \frac{\int \mathbf{1} \{v_{ij} - p_j \geq v_{ik} - p_k \text{ for all } k, \text{ and } v_{ij} - p_j \geq 0\} (c_i - r(w_i)) dF(i)}{\int \mathbf{1} \{v_{ij} - p_j \geq v_{ik} - p_k \text{ for all } k, \text{ and } v_{ij} - p_j \geq 0\} dF(i)},$$

and the total and average spending by the market sponsor are, respectively

$$G^r(p) = \int \mathbf{1} \left\{ \max_j \{v_{ij} - p_j\} \geq 0 \right\} r(w_i) dF(i), \quad g^r(p) = \frac{G^r(p)}{\sum_j q_j(p)}.$$

Under a subsidy design $s(\cdot)$ the enrollment function for plan j is

$$q_j^s(p) = \int \mathbf{1} \{v_{ij} - p_j \geq v_{ik} - p_k \text{ for all } k, \text{ and } v_{ij} - p_j \geq -s(w_i)\} dF(i),$$

which simplifies to $q_j^s(p) = q_j(p - s)$ if $s(\cdot)$ is a (constant) uniform subsidies. Total and average market sponsor spending are $G^s(p) = \int \mathbf{1} \left\{ \max_j \{v_{ij} - p_j\} \geq -s(w_i) \right\} s(w_i) dF(i)$, and $g^s(p) = G^s(p) / \left(\sum_j q_j^s(p) \right)$. Under uniform subsidies $s(w_i) = s$ (for all w_i) these are simply $G^s(p) = s \sum_j q_j(p - s)$, and $g^s(p) = s$.

Proof of Proposition 1. Let p^* be the vector of equilibrium prices when the risk adjustment $r(w_i)$ is adopted and there are no subsidies. Since contracts are homogeneous,

$p_1^* = p_2^* = \dots = p_j^* = \bar{p}$, and all insurers have the same average costs and obtain the same per-enrollee risk adjustment transfer equal to $g^r(p^*)$. In equilibrium $\bar{p} = AC_j^r(p^*) = AC_j(p^*) - g^r(p^*)$.

Consider now the alternative policy in which there is no risk adjustment, while subsidies are $s(w_i) = \bar{s} = g^r(p^*)$, for all w_i . The price vector $\hat{p} = (\hat{p}_1, \hat{p}_2, \dots, \hat{p}_J)$ with $\hat{p}_j = \hat{p} = AC_j(p^*)$ for all j is then the new equilibrium, since $\hat{p} = \bar{p} + \bar{s}$, and thus $\hat{p} = AC_j(p^*) - g^r(p^*) + \bar{s} = AC_j(\hat{p} - \bar{s})$, so insurers break even. At this equilibrium, enrollment is the same for all types since net-of-subsidy prices are the same as in the original equilibrium and the sponsor spending is the same since $g^s(\hat{p}) = \bar{s} = g^r(p^*)$. ■

Proof of Proposition 2. Begin with the equilibrium p^* under the risk adjustment $r(w_i)$. The first order condition in equation (1) must hold, where we replace $AC(\cdot)$ with $AC^r(\cdot)$. Since we assume symmetry of contracts, they all receive the same average risk adjustment transfer equal to $g^r(p^*)$. The first order condition can then be rewritten as

$$p_j^* = \underbrace{AC_j(p^*) - g^r(p^*)}_{AC_j^r(p^*)} - \frac{q_j(p^*)}{\partial q_j(p^*)/\partial p_j} (1 - \partial AC_j^r(p^*)/\partial p_j). \quad (6)$$

Consider now a case with no risk adjustment and a uniform subsidy that is given by

$$s(w_i) = s^* = g^r(p^*) + \frac{q_j(p^*)}{\partial q_j(p^*)/\partial p_j} \frac{\partial g^r(p^*)}{\partial p_j} \quad \text{for all } w_i. \quad (7)$$

This level of subsidy is constructed so that it satisfies two key properties.

First, it gives rise to a (symmetric) equilibrium in which each insurer j sets premium $\hat{p}_j = p_j^* + s^*$. To see this, note that with subsidy s^* and no risk adjustment, equilibrium must satisfy the following first-order condition

$$\hat{p}_j = AC_j(\hat{p} - s^*) - \frac{p_j(\hat{p} - s^*)}{\partial q_j(\hat{p} - s^*)/\partial p_j} (1 - \partial AC_j(\hat{p} - s^*)/\partial p_j). \quad (8)$$

Replacing $p_j^* = \hat{p}_j - s^*$ implies $p_j^* + s^* = AC_j(p^*) - \frac{q_j(p^*)}{\partial q_j(p^*)/\partial p_j} \left(1 - \frac{\partial AC_j(p^*)}{\partial p_j}\right)$, and substituting for s^* its construction from equation (7) yields the original first-order condition from equation

(6) so that the first order condition in equation (8) must hold.

The second property of this particular construction of s^* is that $s^* = g^s(\hat{p}) = g^r(p^*) + \frac{q_j(p^*)}{\partial q_j(p^*)/\partial p_j} \frac{\partial g^r(p^*)}{\partial p_j} < g^r(P^*)$, where the inequality follows from the fact that demand slopes downward — $\frac{q_j(p^*)}{\partial q_j(p^*)/\partial p_j} < 0$ — and regular risk adjustments under adverse selection imply that $\frac{\partial g^r(p^*)}{\partial p_j} > 0$. This concludes the proof because it shows a subsidy design in which demand and insurance allocations remain the same as under risk adjustments, but sponsor expenditure is lower. ■

B Details of counterfactual simulations

Counterfactual policies varying budget. Here we describe in more details how we obtain the results in Section 4, where we solve for equilibria over a large grid of policy parameters and then find the policies that maximize enrollment subject to sponsor budget.

For risk adjustment, we set up a grid of values of $\boldsymbol{\mu} = \{\mu^1, \mu^2, \dots, \mu^N\}$ and for every region-year market we solve for the equilibrium $p^{sol,n}$ under risk adjustment $r(w_i) = \mu^n w_i$, where $sol = pc, ipc$ distinguishes between perfect and imperfect competition.¹⁴ We use the resulting combinations of $\{\mu^n, p^{sol,n}\}$ to compute total enrollment $Q^{sol,n} = \sum_j q_j(p^{sol,n})$ and sponsor total ($G^{sol,n}$) and per-buyer ($g^{sol,n}$) spending. We then expand our grid by interpolating values using local polynomials, and find the value n^* that solves

$$\max_n Q^{sol,n} \text{ s.t. } G^{sol,n} \leq B,$$

where B is the total budget.

The procedure for subsidies is analogous. For uniform subsidies we set up a grid of values $\mathbf{S} = \{s^1, s^2, \dots, s^N\}$, while for targeted subsidies we consider the two-dimensional grid of values $\lambda^k = (\lambda_0^n, \lambda_1^m)$ where $\lambda_0^n = s^1, s^2, \dots, s^N$ while $\lambda_1^m = \{0.05, 0.1, 0.2, 0.5\}$.

¹⁴Under perfect competition (Figure 5a), $p^{pc,n}$ solves $p^{pc,n} = AC^r(p^{pc,n})$. Under imperfect competition we consider two cases. For monopoly (Figure 5) we find the profit-maximizing premium $p^{ipc,n}$, for multiproduct oligopoly (Figures 6-7 and Table 2) we use a nonlinear optimizer to find the vector $p^{ipc,n}$ that minimizes the total square deviations from the Bertrand-Nash first order conditions.

Noisy risk adjustment and income-targeted subsidies. At the end of Section 4.4 we modify our analysis to include “imperfect” risk adjustment and income-based subsidies.

For the former, we construct a noisy signal of risk for every individual in our data. That is, we set $w_i = \exp(\rho_i)/\mathbb{E}\exp(\rho_i) + \omega_i$, where ω_i is drawn independently from a normal distribution with standard deviation that is set to ensure that the R^2 of a linear regression of $\exp(\rho_i)$ on w_i is equal to 0.3, the value that we consider following the literature on risk adjustment.

For the case of income-based subsidies we simply set $w_i = \text{Income}_i/\mathbb{E}(\text{Income}_i)$. We then follow the steps we followed for targeted subsidies, and consider $s(w_i) = \lambda_0(1 - \lambda_1(1 - 1/w_i))$, varying λ_0 and λ_1 over a grid.